vDef-Web: A Case-Study on Building a Science Gateway Around a Research Code

Alexander Dunkel
Department of Mathematics
Louisiana State University
Baton Rouge, LA, USA
fdunke1@lsu.edu

Blaise Bourdin
Department of Mathematics
Louisiana State University
Baton Rouge, LA, USA
bourdin@lsu.edu

Steven R. Brandt
Center for Computation and Technology
Louisiana State University
Baton Rouge, LA, USA
sbrandt@cct.lsu.edu

***Abstract*—Many research codes assume a user's proficiency with high-performance computing tools, which often hinders their adoption by a community of users. Our goal is to create a user-friendly gateway to allow such users to leverage new capabilities brought forward to the fracture mechanics community by the phase-field approach to fracture, implemented in the open source code `vDef`.**

We leveraged popular existing tools for building such frameworks: Agave, Django, and Docker, to build a Science Gateway that allows a user to submit a large number of jobs at once. We use the Agave framework to run jobs and handle all communications with the high-performance computers, as well as data sharing and tracking of provenance. Django was used to create a web application. Docker provided an easily deployable image of the system, simplifying setup by the user.

The result is a system that masks all interactions with the high-performance computing environment and provides a graphical interface that makes sense for scientists. In the common situation of parameter sweeps our gateway also helps the scientists comparing outputs of various computations using a matrix view that links to individual computations.

***Index Terms*—Agave, high-performance computing, graphical user interface, science gateway, parameter sweep, phase-field models, fracture mechanics**

I. Introduction

Variational phase-field models of fracture, introduced in [1], [2] as regularization of Francfort and Marigo's variational approach to fracture [3], have brought a paradigm shift to the study of fracture of brittle solids. Classical methods, born out of Griffith's seminal works [4]–[6] require ad-hoc hypotheses on crack path and have difficulty handling nucleation, merging, or interactions between cracks. In contrast, phase-field models can predict path [7], [8], nucleation [9], branching or merging, in two and three spatial dimensions. These features have been leveraged to study fracture in complex situations such as drying of colloidal suspension [10], thermal cracks in ceramics [11], and fracture of thin films [12]–[14]. They are also well-suited to the study of complex coupled problems where fracture is just one aspect of the physics involved. Such problems include hydraulic fracturing [15], [16], or the understanding of the failure of heterogeneous materials [17]–[19]. The open-source code `vDef` [20], which has been used in most of the references above and deployed in several industry research centers, can potentially benefit an even broader community of scientists and engineers who are not experts in fracture mechanics or high-performance computing. However, as with many research codes, using `vDef` requires familiarity with classical tools of high-performance computing (command line, file transfer, batch execution on remote supercomputing, and post-processing) which is not widespread amongst students and researchers with basic training in engineering and basic science. In the context of parameter sweeps, when a user tries to understand or characterize a class of problem, data management and the need for a synthetic view of large result sets becomes a major issue, even for users familiar with high-performance computing.

In this article, we describe `vDef-Web` [21], a science gateway developed to address both issues. It is built upon:

- the Agave framework to handle all communications with high-performance computing systems including data access, data sharing, job submission, and provenance tracking;
- Django to create the `vDef-Web` application as a graphical user interface for the user; and
- Docker to build an image to deploy `vDef-Web` easily on the user's local computer.

Users are presented with a web interface in which they can describe a problem, upload templates of input files, and specify a set of parameters to explore. Jobs are then automatically deployed, post-processed, and summarized in a compact matrix view from which details of each computation and data files can be retrieved.

II. Methodology

A. The Agave Platform

To access the high-performance computers and their data, we are using the Agave Platform. The Agave platform is

Support for AD and BB was provided in part by U.S. National Science Foundation grant DMS-1716763 under the Applied Mathematics program. Support for SB was provided in part by U.S. National Science Foundation grant OAC-1450437. Some numerical experiments were performed using resources of the Extreme Science and Engineering Discovery Environment (XSEDE), supported by National Science Foundation grant OCI-1053575. Others were conducted with high performance computing resources provided by Louisiana State University (http://www.hpc.lsu.edu).

Presented at Gateways 2019, San Diego, CA, USA, September 23–25, 2019. https://osf.io/meetings/gateways2019/

an open-source science-as-a-service API platform that allows researchers to manage and share data, and run and share code in a reproducible way. We use Agave to submit the jobs to the selected machines. Agave handles all the authentications and interactions with the machines, hiding those interactions from users of our system. Furthermore, Agave allows users to share the registered machines, applications, and results of finished jobs, allowing the users to collaborate efficiently.

B. Django

Django is a free and open-source web framework based on Python. It is easy to use and promotes rapid web development. Django has a Model-View-Template architecture. It consists of a template describing how the data is displayed, a view, i.e., the logic describing which data gets displayed to the user, and an object-relational mapper used by the view to communicate with the database. We use an SQLite database to store data about submitted jobs such as job name, job id, and output data used for visualization. Django includes many security features such as protection against cross-site scripting, cross-site request forgery, clickjacking, and SQL injection.

C. Docker

Docker is an open source software to create, deploy, and run applications by using containers. The container is a closed environment that contains everything that is needed to run the application. Besides the application itself, the container contains all dependencies, libraries, and configuration files that the application needs. This abstraction separates the application from its underlying environment and infrastructure. By containerizing the application, users can quickly deploy and execute `vDef-Web`.

D. Putting it All Together

To start `vDef-Web`, the user runs *docker-compose*. Once the web application is running, a user can log in providing their Agave credentials. The user can then submit new jobs, check the status of jobs, visualize previously submitted jobs, or share results, applications, and systems, all within the web application. The web application interacts with the Agave framework which interacts with the high-performance computing systems. When submitting a job, the user can provide an email address to receive an email alert through Agave's built-in notification service, when the job has finished and is ready to be visualized.

Fig. 1 shows how the pieces of the system interact with one another. The user interacts directly only with `vDef-Web` which uses Agave to hide the direct usage of the `vDef` and the high-performance computing environment.

E. Related Work

This science gateway is similar to a number of other small projects built around a particular community with distinct science needs. Like many such projects, e.g., [22]–[26], it follows the guidance of "Authoring a Science Gateway Cookbook" [27] and leverages Agave [28] to provide basic infrastructure, credential management, file transfer, job management, and provenance.

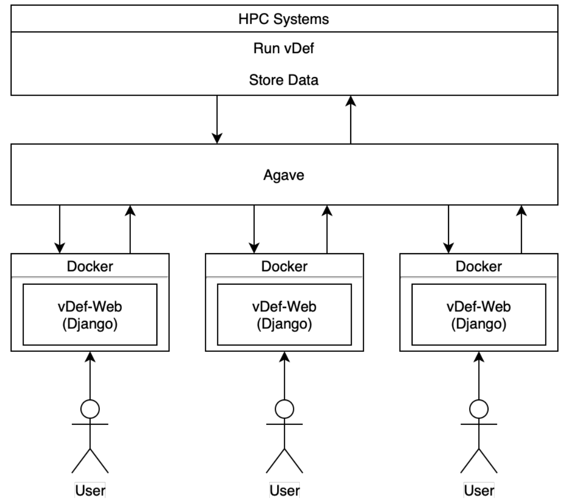

Fig. 1. Schematic of workflow.

Django is likewise, a common ingredient in a number of science gateways [29], [30], etc., as it simplifies integration with a database and creation of websites.

III. Results

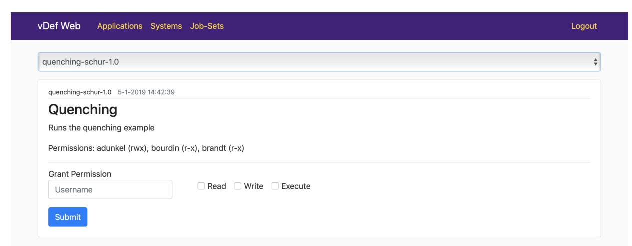

Fig. 2. Screenshot showing the view of an application.

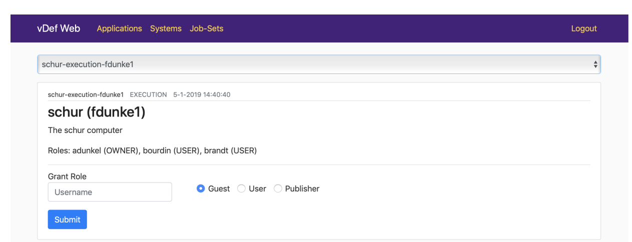

Fig. 3. Screenshot showing the view of a system.

When users are logged in to our gateway, they can view the systems and applications that are available to them and have the option to share them with other users as shown in Fig. 2 and Fig. 3. A user can get access to systems and applications either through another user that shares their systems and applications or by creating new systems and applications using the Agave API, which would require some familiarity with the command-line interface.

If the user selects an application to submit jobs, the user selects the input files required for the job submission. `vDef` requires two input files: a GEO file describing the geometry of the material and a YAML file describing the mechanical properties. The user may wish to perform a parameter sweep over any of the properties that they set within the input files. To let `vDef-Web` know which variable to use for the parameter sweep, the user needs to mark the variable. We mark the variables by two open curly brackets before the variable and two close curly brackets after the variable, i.e., {{variable}}. When an input file is uploaded, the web application scans the file for this expression and gives the option to select a range and number of parameters for the parameter sweep. In addition to the parameter options, the user can select the archive and execution system, and set the number of nodes, the number of processors per node, the maxium run time, and the queue. An optional email can be provided to receive an alert when the jobs have finished. Finally, a unique name for the collection of jobs, which we call a job-set, needs to be selected by the user to identify the job-set. The procedure of submitting a job through the web application can be seen in Fig. 4 and Fig. 5.

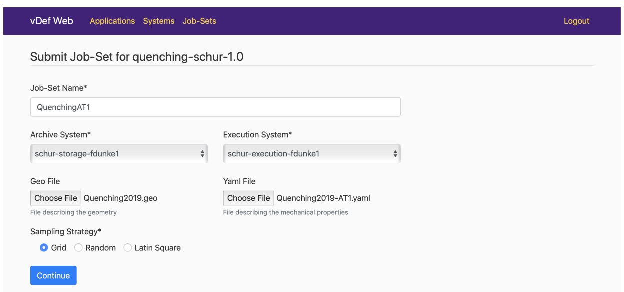

Fig. 4. Screenshot showing the job submission procedure where the user can set the job-set name, select the archive and execution system, and upload the input files.

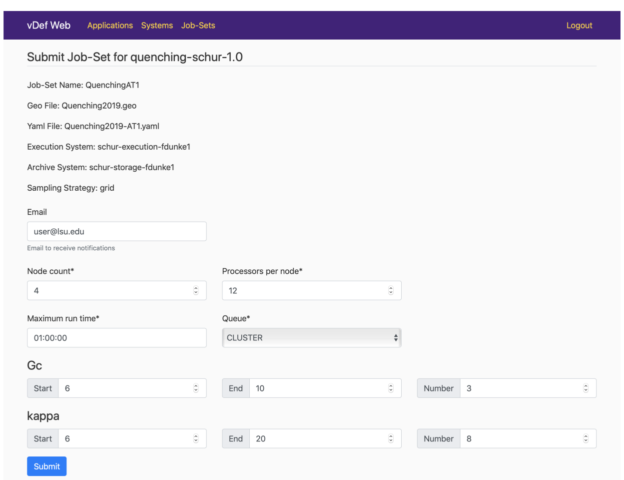

Fig. 5. Screenshot showing the job submission procedure where the user can set an email address, the number of nodes, the number processors per node, a maximum run time, a queue, and the parameter options.

To visualize the results of a job-set, the user can select the specific job-set from a list of all submitted sets. When the user selects a set to visualize, `vDef-Web` retrieves the current status of each job. If the job has finished, the gateway downloads and stores the resulting image. We use a plot of points for the visualization, where each point represents one computation. The color of the point indicates the status of the job: A blue point is a finished job, an orange point is a job that is running, and a red point is a job that has failed. When the user hovers over one of the points for a job that has finished, the system displays the resulting image. When clicking on one of the jobs, the system displays the corresponding output files which can be downloaded by the user on demand. Fig. 6 shows the visualization of a job-set for the numerical simulation of Yuse and Sano's classical glass quenching experiment [31], [32].

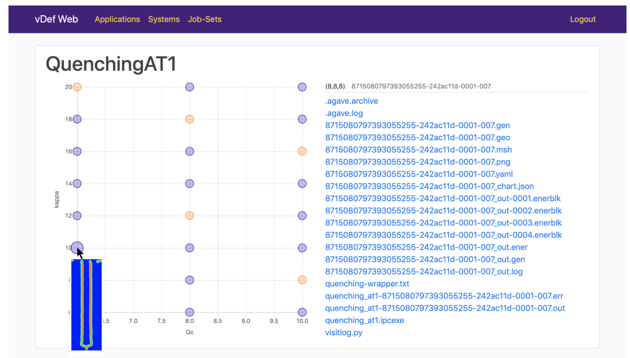

Fig. 6. Screenshot showing the visualization of a job-set.

In this experiment, a slab of glass is heated up then quenched at a constant speed in a bath of cold water. Whereas the experimental and material parameters are many, a simple dimensional analysis can be used to show that the entire parameter space can be described in terms of two non-dimensional parameters, quenching speed V and temperature contrast $\Delta\theta$, or fracture toughness G_c and thermal diffusivity κ. In experiments, three broad classes of outcomes are ob-

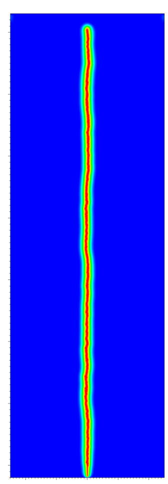

Fig. 7. Quenching Result for $G_c = 10$ and $\kappa = 20$.

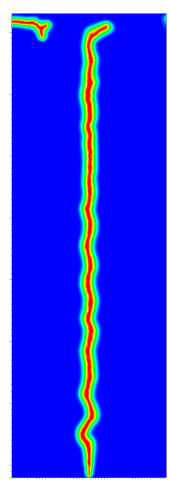

Fig. 8. Quenching Result for $G_c = 6$ and $\kappa = 16$.

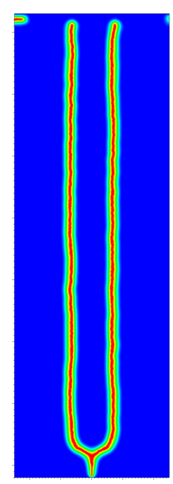

Fig. 9. Quenching Result for $G_c = 6$ and $\kappa = 14$.

served: propagation of a simple straight crack, of a single crack oscillating along a sine-like path, or "erratic" regime in which complex crack paths, and possibly branching is observed.

A parameter sweep for this problem can be used to build a "phase diagram," *i.e.* identify the regions of parameter space where behavior is attained, and compare it to the experimental literature in order to validate algorithms and models.

We performed these computations for different values for the fracture toughness (G_c) and thermal conductivity (κ) of the glass. By hovering over the computations, the researcher can quickly identify the different fracture patterns. In Fig. 7-9, we can see different fracture patterns for different values of G_c and κ.

IV. Conclusion

The phase-field approach to fracture has brought great advances to the community of fracture mechanics. With `vDef`, the numerical implementation of this approach, the community can benefit significantly. The nature of scientific codes requires knowledge of high-performance computing tools and the command-line interface. The lack of familiarity can be a hurdle to the scientists, which is usually only overcome with difficulty. Our gateway provides a straightforward graphical interface that researchers in the fracture mechanics community can use to run simulations. Furthermore, they can perform parameter sweeps effortlessly and analyze the results in a matrix view.

References

[1] B. Bourdin, G. Francfort, and J.-J. Marigo, "Numerical experiments in revisited brittle fracture," *J. Mech. Phys. Solids*, vol. 48, no. 4, pp. 797–826, 2000.

[2] B. Bourdin, G. Francfort, and J.-J. Marigo, "The variational approach to fracture," *J. Elasticity*, vol. 91, no. 1-3, pp. 1–148, 2008.

[3] G. Francfort and J.-J. Marigo, "Revisiting brittle fracture as an energy minimization problem," *J. Mech. Phys. Solids.*, vol. 46, no. 8, pp. 1319–1342, 1998.

[4] A. A. Grith, "The phenomena of rupture and ow in solids," *Philos. Trans. Roy. Soc. London*, vol. 221, 1921.

[5] G. Irwin, "Analysis of stresses and strains near the end of a crack traversing a plate," *J. Appl. Mech.*, vol. 24, pp. 361–364, 1957.

[6] J. Rice, "A path independent integral and the approximate analysis of strain concentration by notches and cracks," *J. Appl. Mech.*, vol. 35, pp. 379–386, 1968.

[7] A. Mesgarnejad, B. Bourdin, and M. M. Khonsari, "Validation simulations for the variational approach to fracture," *Comp. Methods Appl. Mech. Engng.*, vol. 290, pp. 420–437, 2015.

[8] K. H. Pham, K. Ravi-Chandar, and C. M. Landis, "Experimental validation of a phase-field model for fracture," *International Journal of Fracture*, vol. 205, no. 1, pp. 83–101, 2017.

[9] E. Tanné, T. Li, B. Bourdin, J.-J. Marigo, and C. Maurini, "Crack nucleation in variational phase-field models of brittle fracture," *J. Mech. Phys. Solids*, vol. 110, pp. 80–99, 2018.

[10] C. Maurini, B. Bourdin, G. Gauthier, and V. Lazarus, "Crack patterns obtained by unidirectional drying of a colloidal suspension in a capillary tube: experiments and numerical simulations using a two-dimensional variational approach," *Int. J. Fracture*, vol. 184, no. 1-2, pp. 75–91, 2013.

[11] B. Bourdin, J.-J. Marigo, C. Maurini, and P. Sicsic, "Morphogenesis and propagation of complex cracks induced by thermal shocks," *Phys. Rev. Lett.*, vol. 112, no. 1, p. 014301, 2014.

[12] A. León-Baldelli, J.-F. Babadjian, B. Bourdin, D. Henao, and C. Maurini, "A variational model for fracture and delamination of thin films under in-plane loading," *J. Mech. Phys. Solids*, vol. 70, pp. 15–32, 2014.

[13] A. Mesgarnejad, B. Bourdin, and M. M. Khonsari, "A variational approach to the fracture of brittle thin films under out of plane loading," *J. Mech. Phys Solids*, vol. 61, no. 11, pp. 2360–2379, 2013.

[14] A. León-Baldelli, B. Bourdin, J.-J. Marigo, and C. Maurini, "Fracture and debonding of a thin film on a stiff substrate: analytical and numerical solutions of a one-dimensional variational model," *Cont. Mech. Thermodyn.*, vol. 25, no. 2-4, pp. 243–268, 2013.

[15] K. Yoshioka and B. Bourdin, "A variational hydraulic fracturing model coupled to a reservoir simulator," *Int. J. Rock Mech. Min.*, vol. 88, pp. 137 – 150, 2016.

[16] C. Chukwudozie, B. Bourdin, and K. Yoshioka, "A variational phase-field model for hydraulic fracturing in porous media," *Comp. Meth. Appl. Mech. Engng.*, vol. 347, pp. 957–982, 2019.

[17] S. Brach, M. Hossain, B. Bourdin, and K. Bhattacharya, "Effective toughness of heterogeneous media: Anisotropy effects." Under review, 2019.

[18] C.-J. Hsueh, L. Avellar, B. Bourdin, G. Ravichandran, and K. Bhattacharya, "Stress fluctuation, crack renucleation and toughening in layered materials," *J. Mech. Phys Solids*, vol. 120, pp. 68 – 78, 2018. Special issue in honor of Ares J. Rosakis on the occasion of his 60th birthday.

[19] M. Hossain, C.-J. Hsueh, B. Bourdin, and K. Bhattacharya, "Effective toughness of heterogeneous media," *J. Mech. Phys. Solids*, vol. 71, pp. 320–348, 2014.

[20] B. Bourdin, "vDef: a reference implementation of variational phase-field models of fracture." https://doi.org/10.5281/zenodo.3242132, 2014.

[21] "vDef-Web." https://doi.org/10.5281/zenodo.3402530, 2019.

[22] S. Yuan, Q. J. Chen, R. Dooley, S. R. Brandt, and L. Zhu, "A sustainable collaboratory for coastal resilience research," Sep 2018.

[23] M. R. Hanlon, M. Vaughn, S. Mock, R. Dooley, W. Moreira, J. Stubbs, C. Town, J. Miller, V. Krishnakumar, E. Ferlanti, *et al.*, "Araport: an application platform for data discovery," *Concurrency and Computation: Practice and Experience*, vol. 27, no. 16, pp. 4412–4422, 2015.

[24] M. R. Hanlon, M. Vaughn, S. Mock, R. Dooley, W. Moreira, J. Stubbs, C. Town, J. Miller, V. Krishnakumar, E. Ferlanti, *et al.*, "The arabidopsis information portal: an application platform for data discovery," in *2014 9th Gateway Computing Environments Workshop*, pp. 38–41, IEEE, 2014.

[25] C. M. Lushbough, E. Z. Gnimpieba, and R. Dooley, "Life science data analysis workflow development using the bioextract server leveraging the iplant collaborative cyberinfrastructure," *Concurrency and Computation: Practice and Experience*, vol. 27, no. 2, pp. 408–419, 2015.

[26] B. D. Corrie, N. Marthandan, B. Zimonja, J. Jaglale, Y. Zhou, E. Barr, N. Knoetze, F. M. Breden, S. Christley, J. K. Scott, *et al.*, "ireceptor: A platform for querying and analyzing antibody/b-cell and t-cell receptor repertoire data across federated repositories," *Immunological reviews*, vol. 284, no. 1, pp. 24–41, 2018.

[27] S. Marru, R. Dooley, N. Wilkins-Diehr, M. Pierce, M. Miller, S. Pamidighantam, and J. Wernert, "Authoring a science gateway cookbook," in *2013 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 1–3, IEEE, 2013.

[28] R. Dooley, S. R. Brandt, and J. Fonner, "The agave platform: An open, science-as-a-service platform for digital science," in *Proceedings of the Practice and Experience on Advanced Research Computing*, p. 28, ACM, 2018.

[29] M. A. Miller, W. Pfeiffer, and T. Schwartz, "The cipres science gateway: enabling high-impact science for phylogenetics researchers with limited resources," in *Proceedings of the 1st Conference of the Extreme Science and Engineering Discovery Environment: Bridging from the eXtreme to the campus and beyond*, p. 39, ACM, 2012.

[30] S. Cholia, D. Skinner, and J. Boverhof, "Newt: A restful service for building high performance computing web applications," in *2010 Gateway Computing Environments Workshop (GCE)*, pp. 1–11, IEEE, 2010.

[31] A. Yuse and A. Sano, "Transition between crack patterns in quenched glass plates," *Nature*, vol. 362, pp. 329–330, 1993.

[32] A. Yuse and M. Sano, "Instability of quasi-static crack patterns in quenched glass plates," *Physica D*, vol. 108, pp. 365–378, 1997.